

PATENT

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In Re Application Of: Hawkes, et al.

Group No.:

Serial No.: To Be Assigned

Docket No. 30018432-2

Filed: October 28, 2003

Confirmation No.:

For: **Simulation At Two Different Levels of Complexity**

CLAIM OF PRIORITY TO AND
SUBMISSION OF CERTIFIED COPY OF UNITED KINGDOM APPLICATION
PURSUANT TO 35 U.S.C. §119

Commissioner for Patents
P.O. Box 1450
Alexandria, Virginia 22313-1450

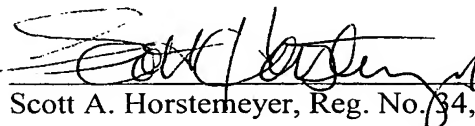
Sir:

In regard to the above-identified pending patent application and in accordance with 35 U.S.C. §119, Applicant hereby claims priority to and the benefit of the filing date of United Kingdom patent application entitled, "Simulation at Two Different Levels of Complexity", filed April 30, 2003, and assigned serial number 0309958.7. Further pursuant to 35 U.S.C. §119, enclosed is a certified copy of the United Kingdom patent application

Respectfully Submitted,

**THOMAS, KAYDEN, HORSTEMEYER
& RISLEY, L.L.P.**

By:

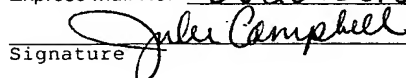

Scott A. Horstemeyer, Reg. No. 34,183

100 Galleria Parkway, Suite 1750
Atlanta, Georgia 30339
770-933-9500

EXPRESS MAIL

I hereby certify that this correspondence is being deposited with the United States Postal Service as "Express Mail Post Office to Addressee" in an envelope addressed to: Commissioner for Patent, Mail Stop Patent Application, P.O. Box 1450, Alexandria, VA, on

Oct 28, 2003
Express Mail No. EV 26933152945

Signature 




The Patent Office
Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ

I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.

Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.



Signed

Dated

22 October 2003

Patents Form 1/77

Patent Act 1977
(Rule 13)



THE PATENT OFFICE
E
30 APR 2003

01MAY03 E804040-1 DD1463
POL/7700 0.00-0309958.7

Request for grant of a patent

(See the notes on the back of this form. You can also get an explanatory leaflet from the Patent Office to help you fill in this form)

The Patent Office

Cardiff Road
Newport
South Wales
NP10 8QQ

1. Your reference	30018432-1 GB		
2. Patent application number (The Patent Office will fill in this part)	0309958.7		30 APR 2003
3. Full name, address and postcode of the or of each applicant (underline all surnames)	Hewlett-Packard Development Company, L.P. 20555 S.H. 249 Houston, TX 77070 USA		
Patents ADP number (if you know it)	855 78 86 00 1		
If the applicant is a corporate body, give the country/state of its incorporation	TEXAS		
4. Title of the invention	Simulation at Two Different Levels of Complexity		
5. Name of your agent (if you have one)	Richard A. Lawrence Hewlett-Packard Ltd, IP Section Filton Road, Stoke Gifford Bristol BS34 8QZ		
"Address for service" in the United Kingdom to which all correspondence should be sent (including the postcode)			
Patents ADP number (if you know it)	-156308 3001		
6. If you are declaring priority from one or more earlier patent applications, give the country and the date of filing of the or of each of these earlier applications and (if you know it) the or each application number	Country	Priority application number (if you know it)	Date of filing (day / month / year)
7. If this application is divided or otherwise derived from an earlier UK application, give the number and the filing date of the earlier application	Number of earlier application		Date of filing (day / month / year)
8. Is a statement of inventorship and of right to grant of a patent required in support of this request? (Answer 'Yes' if: a) any applicant named in part 3 is not an inventor, or b) there is an inventor who is not named as an applicant, or c) any named applicant is a corporate body. See note (d))	Yes		

Patents Form 1/77

9. Enter the number of sheets for any of the following items you are filing with this form. Do not count copies of the same document

Continuation sheets of this form

Description

21

Claim(s)

4

Abstract

1

Drawing(s)

3+3

10. If you are also filing any of the following, state how many against each item.

Priority documents

-

Translations of priority documents

-

Statement of inventorship and right to grant of a patent (Patents Form 7/77)

-

Request for preliminary examination and search (Patents Form 9/77)

1

Request for substantive examination (Patents Form 10/77)

-

Any other documents (please specify)

Fee Sheet

11.

I/We request the grant of a patent on the basis of this application.

Signature

Richard A. Lawrence

Date

30/4/03

12. Name and daytime telephone number of person to contact in the United Kingdom

Meg Joyce Tel: 0117-312-9068

Warning

After an application for a patent has been filed, the Comptroller of the Patent Office will consider whether publication or communication of the invention should be prohibited or restricted under Section 22 of the Patents Act 1977. You will be informed if it is necessary to prohibit or restrict your invention in this way. Furthermore, if you live in the United Kingdom, Section 23 of the Patents Act 1977 stops you from applying for a patent abroad without first getting written permission from the Patent Office unless an application has been filed at least 6 weeks beforehand in the United Kingdom for a patent for the same invention and either no direction prohibiting publication or communication has been given, or any such direction has been revoked.

Notes

- If you need help to fill in this form or you have any questions, please contact the Patent Office on 08459 500505.
- Write your answers in capital letters using black ink or you may type them.
- If there is not enough space for all the relevant details on any part of this form, please continue on a separate sheet of paper and write "see continuation sheet" in the relevant part(s). Any continuation sheet should be attached to this form.
- If you have answered 'Yes' Patents Form 7/77 will need to be filed.
- Once you have filled in the form you must remember to sign and date it.
- For details of the fee and ways to pay please contact the Patent Office.

SIMULATION AT TWO DIFFERENT LEVELS OF COMPLEXITY

Field of the Invention

5 The present invention relates to improvements in simulation. It is particularly applicable to the simulation of creatures for use in two different complexities of simulation.

10

Background of the Invention

As will be indicated later, although the present invention has broader application, it arose in connection with the
15 simulation of creatures. This area of art will now be discussed.

One of the problems with studying life sciences is that changes in ecosystems and the evolution of life forms
20 takes place over long periods of time. Consequently, life processes have been simulated in a variety of formats, thus allowing the study of the interaction of creatures and the emerging population patterns.

25 An animat is an artificial creature which exhibits an active and purposeful behaviour in an environment which presents it with problems - obstacles, navigation, hazards, goals. Animats are often inspired by biological organisms, which are observed to solve complex problems in
30 real physical environments. Indeed, animats may be designed to be as close as possible to the original biological organisms, so as to model the potential behaviour and/or evolution of such organisms.

It will be understood that use of the word "animat" is intended to mean not only artificial creatures that are directly inspired by real biological creatures, but also
5 imaginary animals and other autonomous agents such as real or imagined robots or "cyborgs" (part animal, part machine).

In order to make such simulations as accurate as possible,
10 it is desirable to lavish as much computational resource on each artificial creature as is available. The method of simulating an animat depends on the available computational resources. If a significant amount of processing power is available, it is possible to perform a
15 high-fidelity simulation (Hi-Fi Sim). Such a high-fidelity simulation will include the use of computationally expensive techniques for behaviour selection, sensing and interaction with the environment and other animats, e.g. use of a form of 3D rigid-body
20 physics simulation to model the environment and the various parts of the body of the animat; use of ray-tracing computer graphics to synthesize visual sensory input; the simulation of the transmission, absorption, and reflection of sound waves, etc.

25

If only a relatively small amount of computational power is available, then a relatively low-fidelity simulation (Lo-Fi Sim) would be utilised to ensure that the simulation progresses at a reasonable rate. Such a Lo-Fi
30 Sim will use computationally less expensive techniques, such as only using a two dimensional (2D) model of the environment, and employing a simplified model of the body of the animat e.g. treating it as a point mass with

predetermined locomotion characteristics (speed, rate of turn etc).

It is desirable that any predetermined type of animat will
5 function in a consistent manner regardless of the level of simulation.

Statement of the Invention

10

In a first aspect, the present invention provides a method of simulating a creature for use in two different complexities of simulation, the method comprising utilising a model of the creature that comprises at least
15 two portions, a first portion which contains functions for use in both of said different complexities of simulation; and a second portion comprising two alternative versions, a first version for use in one of said different complexities of simulation, and a second version for use
20 in the other of said different complexities of simulation.

Preferably, said first portion comprises a behaviour selection mechanism arranged to select the behaviour of said creature.

25

Preferably, said behaviour selection mechanism is arranged to select the behaviour based upon at least one of: the current behavioural state; an energy level of the creature; the environment surrounding the creature; one or
30 more sensory inputs to said creature; an internal state of said creature.

Preferably, said behaviour selection mechanism consists of a set of mutually exclusive behavioural states.

Preferably, the second version is for use in the less
5 complex of the simulations, and is arranged to approximate the functionality of the first version.

Preferably, the first version utilises a neural network.

10 Preferably, said second portion is arranged to execute the selected behaviour.

Preferably, the first version utilises a three dimensional physical simulation of the animat, and the second version
15 utilises a parameterised model of the animat to approximate movement.

In another aspect, the present invention provides a method of simulating the activities of a plurality of creatures,
20 the method comprising utilising at least two modes of simulation: a first mode arranged to simulate the activities of all (or very many) of said creatures; and a second mode arranged to simulate the activity of at least one of said creatures at a more detailed level than said
25 first mode; wherein the model of a creature simulated in both modes of simulation comprises at least two portions: a first portion which contains functions arranged for use in both of said modes of simulation; and a second portion comprising two alternative versions, a first version for
30 use in said first mode of simulation, and a second version for use in the second mode.

In a further aspect, the present invention provides a computer program arranged to perform the method.

In another aspect, the present invention provides a recordable medium comprising a computer program.

In a further aspect, the present invention provides a simulator device arranged to simulate a creature in two different complexities of simulation; the device being arranged to utilise a model of the creature that comprises at least two portions; a first portion which contains functions used in both of said different complexities of simulation; and a second portion comprising two alternative versions, a first version used in one of said different complexities of simulation, and second version used in the other of said different complexities of simulation.

In a further aspect, the invention provides a method of simulating a process at two different levels of complexity, the method comprising utilising a model that comprises at least two portions, a first portion which contains functions for use in both of said different complexities of simulation; and a second portion comprising two alternative versions, a first version for use in one of said different complexities of simulation, and a second version for use in the other of said different complexities of simulation.

It is thus appreciated that the present invention has broader application than to the simulation of creatures.

Brief Description of Drawings

For a better understanding of the invention, and to show
5 how embodiments of the same may be carried into effect,
reference will now be made, by way of example, to the
accompanying diagrammatic drawings in which:

Figure 1 shows a block diagram illustrating the simulation
components that are independent of, and the components
10 that are specific to, the different complexities of
simulation, according to a preferred embodiment of the
present invention;

Figure 2 illustrates the different modes and associated
transition probabilities between the modes of a high level
15 brain, according to the preferred embodiment of the
present invention; and

Figures 3A and 3B illustrate the generic application of
aspects of the invention to simulation.

20

Detailed Description of Embodiments of the Invention

Firstly, a preferred embodiment of the invention for use
in the simulation of creatures will be described. After
25 this, a further, generic, embodiment of the invention will
be discussed, together with its application to simulation
problems.

In the preferred embodiment of the present invention, a
30 single computational resource is utilised to simulate a
relatively large animat population (e.g. 100,000's) in a
Lo-Fi Sim, with an associated Hi-Fi Sim acting as a slave
program to the Lo-Fi Sim.

The nature of the computational resource will not be discussed in detail here, as conventional computational resources can be employed, as will be well understood by the skilled person. Any computational system with
5 appropriate processor, memory and (where required) means of receiving instructions and displaying or otherwise outputting data can be provided. Simultaneous simulation of a large number of entities is a computationally intense task, so real time simulation, for example, is likely to
10 require reasonably high computational performance. As will be understood, however, aspects of the present invention allow for most effective use of whatever computational resource is available for conducting the simulation.

15

The Hi-Fi Sim is engaged by the Lo-Fi Sim when a more detailed simulation of the activities of one or more animats is required e.g. when an animat from that population is selected for closer inspection by a user, or
20 alternatively when an animat requires more detailed simulation. For example, this could occur when the animat encounters another animat and/or undergoes a change in its environment (e.g. encounters water or food) that could lead to an observable change in the behaviour of the
25 animat.

It is desirable that there should be a high correspondence between the detailed model of the Hi-Fi Sim and more abstract representations used in the Lo-Fi Sim. For
30 instance, it would clearly be undesirable if the Lo-Fi sim were to simulate a particular animat as being attracted to another animat, whilst in the same situation the Hi-Fi Sim were to simulate the particular animat as being repulsed

by this further animat. As the Hi-Fi Sim would be utilised to model the detailed interaction of these two animats, then this would lead to a cyclic behaviour, with the Lo-Fi Sim invoking the Hi-Fi Sim when the animats are close enough to interact, the Hi-Fi Sim then indicating that the animats would repulse each other, the Lo-Fi Sim indicating the animats would attract each other and invoking the Hi-Fi Sim again etc.

- 10 To overcome such problems, in the preferred embodiment the brain of each animat is separated into two parts: a high-level brain (HLB) and a low-level brain (LLB).

The HLB is a behaviour selection mechanism, that is arranged to use external (sensory) events and situations, as well as internal (e.g. simulated bio-chemical and energetic reaction levels) events and situations to decide what behaviour is appropriate at any given time. The HLB of an animat is used in both the Hi-Fi Sim and the Lo-Fi Sim. The internal operation of the HLB is identical in both simulations.

Different LLB's are designed for the Hi-Fi Sim and the Lo-Fi Sim. The LLB is capable of executing the currently selected HLB behaviour. In this particular embodiment, the Hi-Fi Sim LLB uses more computationally expensive techniques to execute the HLB behaviour. The Lo-Fi Sim LLB uses cheaper techniques that are designed to approximate the effect of the Hi-Fi Sim LLB as close as possible.

For example, the Hi-Fi Sim in this embodiment is arranged to use a form of three dimensional (3D) modelled

environment, including 3D physical simulations of each
 animat. The Hi-Fi Sim LLB is consequently arranged to use
 neural networks to control the articulation of the limbs
 of the body of the animat, with the interaction of limbs
 5 with the environment being modelled using rigid-body
 dynamics for physical simulation.

On the other hand, the Lo-Fi Sim has only a 2D (two
 dimensional) model of the environment, with the
 10 corresponding LLB simply using a parameterised model of
 the Hi-Fi Sim locomotion to approximate the movement. For
 instance, the LLB treats the body of the animal as a point
 mass with predetermined parameters such as speed, rate of
 turn etc.

15 By separating the model of the animat into two portions,
 it is possible to ensure consistent animat behaviour in
 both Hi-Fi Sim and Lo-Fi Sim, while using the most
 appropriate techniques within each mode of simulation as
 20 required

Figure 1 shows the components of the simulation that are
 specific to the Hi-Fi Sim ("High Fidelity") and the Lo-Fi
 Sim ("Low Fidelity"). The dotted line 10 denotes the
 25 boundary between the simulation independent and the
 simulation specific components.

As can be seen, in the preferred embodiment the
 environment 20 is split into two portions, the LoFi
 30 environment abstraction 120 and the HiFi environment
 implementation 220. There are correspondingly two
 representations of an animat: the LoFi Animat 30 and the
 HiFi Animat 40.

In both representations there are sensors 50 and actuators 60. In the LoFi representation the sensors are modelled by a computationally inexpensive abstraction 150 and the
5 actuators are modelled by a similarly simplified abstraction 160. In the HiFi representation the sensors are implemented using an accurate and realistic (but computationally expensive) implementation 250 and similarly the actuators are implemented using an accurate
10 and realistic (but computationally expensive) implementation 260.

The sensors and actuators each require an appropriate representation of the simulated environment, and so in the
15 LoFi Sim the process 1210 of computing what the animat senses in the environment 20 is achieved by use of an abstract representation of the physics of sensing 121. Similarly the process 1220 of effecting the changes in the animat's relationship to its environment 20 as a
20 consequence of its current actuator settings 160 is achieved by use of an appropriate abstract representation of the physics of actuation 122.

In an analogous fashion for an instance of a HiFi Sim
25 animat, the process 2210 of determining what the animat senses relies on the HiFi Sim sensors 250 having available a high fidelity implementation of the relevant sensor-physics 221. Similarly the process 2220 of determining the effects of the HiFi Sim actuators 260 requires a high-
30 fidelity model of the physics of actuation 222 in the environment 20.

It should be understood that the word "physics" is used here to refer to the nature of the environment's "laws of physics" and can refer both to simulations of physics systems similar to those familiar from common knowledge of the laws of physics in the real known universe, and also to fictional and imaginary physics systems unlike those known to exist in the real universe. An example of an imaginary physics system is the "laws of physics" of many humorous animated cartoons, where an animated character that runs at high speed off the edge of a cliff will typically continue to run horizontally through the air until it notices that it is in mid-air, at which point it ceases to travel horizontally and instead starts to fall vertically downwards.

Referring again to Figure 1, in an example of the preferred embodiment of the invention the intended application may be a computational simulation of wild animals on the African plains. For each species of animal that is simulated, there will be a LoFi Sim representation and a corresponding HiFi Sim representation.

In this example animats are either simulated antelope or simulated lions. Each individual animal in the simulation will have its own LoFi Animat representation that is computationally inexpensive to calculate updates for and that requires comparatively little storage space (RAM-space or disk-space). A HiFi Animat representation of an individual animal will only be created when necessary for the purposes of the simulation, e.g. when the creature fights or pursues a prey or evades a predator.

In this example of the preferred embodiment, there may be many thousands of simulated antelopes, and many hundreds of lions, each of which is represented in the LoFi Environment 120 as a point with co-ordinates (x,y) on a 2-dimensional (2-d) plane. The LoFi Actuators 160 may simply apply a 2-d vector to the animat's co-ordinates, updating its position. The vector's direction and magnitude would vary as the animat runs, eats, or sleeps. The antelope-animat's LoFi sensors 150 may simply examine the LoFi Sensor Physics representation of the environment 121 to determine the distance and direction of the nearest food source, the nearest water source, the nearest other antelope and the nearest lion. Sources of food and water may simply be represented in the LoFi Environment Abstraction 120 as circular patches on the ground-plane, and so can be represented by three numbers: the (x,y) co-ordinates of the centre of the patch, and its radius. Thus it can be seen that the LoFiSim representation of many thousands of antelope, many hundreds of lions, and many patches of food and many patches of water can be represented without the need for much computation-power or storage-space.

In this example it could be the case that if a lion is sufficiently close to an antelope, if the lion is hungry and if it has enough energy then it will chase the antelope. At this point the LoFiSim abstraction 120 and animat representations 30 are not sufficient to capture the details of the interaction between the two creatures. Consequently a HiFiSim representation of each animal is created, with HiFi sensors 250 and HiFi actuators 260 interacting with the HiFi Environment Implementation 220.

For greater accuracy and realism, the HiFi Environment 220 may involve a 3-d model of the animals' environment, including representations of natural terrain features such as hills and trees and rivers that form potential obstacles or impediments. The HiFi Sensor Physics 221 may use 3D graphics techniques to model directional vision in each HiFi animat, in which case the Sensor HiFi Implementaion would have to model the "retina" of each animat as an array of pixels with appropriate modelling of the animal's optics (field of view, acuity of vision, ect). Similarly the HiFi Actuator Physics 222 may use advanced mathematical modelling of Newtonian kinematics and dynamics such as those known in the prior art "physics engine" commercial middleware software products "Mathengine" (<http://www.mathengine.com>) and "Havok" (<http://www.havok.com>), in which case the Actuators HiFi Implementation 260 would represent the animal's body skeletal frame and musculature in a manner consistent with the physics engine chosen for the HiFi Actuator Physics 222.

It will be appreciated that the HiFi Actuator Physics 222 and the HiFi Sensor Physics 221 are more tightly coupled within the HiFi Environment Implementation 220 than are the corresponding aspects of the LoFi Environment Abstraction 120. Consider (for example) that what the creature senses using vision is affected by the current orientation of the creature's eyes and head, which is in turn determined by the actions of the creature's gaze-actuators (eyeball and neck muscles). It is typically necessary to represent such sensory-motor coupling in High Fidelity simulations but it can be abstracted away from in Low Fidelity representations.

It will furthermore be appreciated that the coupling of 221 and 222 within 220 can extend to the creature's "sensation of self", e.g. modelling an animal's
5 interoception (e.g. sensation of pain/damage, satiation, etc) or the analogous monitoring functions in machines (e.g. monitors/sensors of oil-levels, engine temperatures, tyre pressures, etc).

10 In the example being described here, the HiFi Animat instances of the chasing lion and the chased antelope will clearly require comparatively large amounts of data-storage space and computer-processor calculations to resolve the outcome of the interaction (i.e. whether the
15 antelope escapes or not). By only creating these HiFi Sim representations when they are needed, and destroying them as soon as they are no longer needed, it is possible to give the impression of highly complex high-fidelity simulations of many thousands of creatures while using a
20 computer processor and memory storage system that is actually capable of accurately computing and storing such high-fidelity simulations for only a very small number of such simulated creatures. The HiFiSim Animat representations 40 are only created when necessary, and
25 are only maintained for as long as is necessary. So long as only a few animats are interacting with their environments (and the other animats) in ways requiring accurate and complex computations, the remainder of the creatures in the simulation can be satisfactorily
30 represented using the LoFiSim Animat representation 30.

The LoFi Animat Representation 30 can be seen to include the High Level Brain or HLB 70 which is responsible for

selecting particular behavioural states or modes of action (such as resting, sleeping, exploring, or eating). The transitions between these states are caused by changes in the creature's Internal State 90 (such as energy levels
 5 altering) and by changchanging the abstract sense-data provided by the LoFi Sensors Abstraction 150.

The HiFi Animat Representation 40 includes the HLB 70 and Internal State values 90 but additionally includes a Low
 10 Level Brain or LLB 80 which implements specific behaviours at the level of triggering actuators in a fashion appropriate for each behaviour. The actions of the LLB can be affected or modulated by the Internal State values 90 and by sensory information received from the Sensors HiFi
 15 Implementation 250. The actions of the HLB 70 can also affect the LLB 80. In each instance of the HiFi Animat Representation, the HLB may be affected by the Abstract LoFi sensory data coming from 150, but in this case the data source for 150 is the HiFi to LoFi converter 100.

20 Thus it can be appreciated that in this embodiment the Sensors LoFi Abstraction 150 and the HLB 70 are used in both the LoFi and HiFi complexity-levels of the simulation, while other components are specific solely to
 25 either the LoFi or the HiFi complexity-levels.

In the preferred embodiment, the HLB 70 is a simple stochastic action selection mechanism used in both the Hi-Fi Sim and the Lo-Fi Sim. The underlying mechanism is a
 30 Markov net, a finite-state process where the different states represent different behavioural modes or activities. The probability of transitions from one state to another are dependent on the animat's internal state

variables (i.e. simulated bio-chemical and energetic levels e.g. indicating the energy level or level of fear of the creature) and any sensory input it receives.

- 5 Figure 2 shows a state machine representing three possible states (rest 410, sleep 420 and explore 430) that are used in a simplistic behaviour selection mechanism of an artificial creature.
- 10 The arrows joining each of the states represent the probability of the transition between each state e.g. $RS=0.3$ indicates that the probability of the creature moving from the rest state to the sleep state is 0.3 for any given time interval, $RE=0.5$ indicates that the probability of the creature moving from the rest to the explore state is 0.5 in the same time interval, and $RR=0.2$ indicates that the probability of the creature remaining in the rest state is 0.2 in the same time interval.
- 15
- 20 At any one time, it is assumed that the creature is only ever in one of the possible states. Consequently, whenever the behaviour is reevaluated, only the probability of the behaviour changing from that state (either back to the same state, or to another valid state) is evaluated.
- 25 The value of the probability of a transition from any given state (and either back to the same state or to another state) is between 0 and 1, with a value of 0 meaning there is no chance of transitioning between the specified states, and an output of 1 meaning that this is the only transition that is possible. Clearly, the sum of all of the transition probabilities from a given state is always equal to 1.
- 30

In this preferred embodiment, the creature perceives its environment via sensors. Preferably, such sensors can identify information about the environment (which may be a simulated real or imaginary physical environment). The sensors detect, for example, the number of visible objects, location, object type etc, with processing occurring so as to determine whether the sensed items are likely to affect the behaviour of the creature (has food been sensed?, a rival?, a potential mate?).

10

Additionally, the creature has other inputs, corresponding to internal drives, such as hunger, fatigue, thirst and desire to reproduce.

15 The transition probability between each behavioural state is not fixed but is instead a function dependent upon factors such as sensory input, internal drives and the temperament of the creature.

20 The transition probability values shown in figure 2 are provided by way of example only, and would in fact vary over time depending upon these other factors.

It will be appreciated that the above embodiment is provided by way of example only, and that various other embodiments will be apparent to the skilled person. For instance, an embodiment of the invention might not be applied to all of the creatures within a simulation, but to only a portion of the creatures.

30

It will be appreciated that the simple state-transition network described here is only one possible embodiment of the HLB behaviour-selection mechanism, and other methods

could be used instead. Examples of other methods include those described by:

- T. Tyrrell. *Computational Methods of Action Selection*. PhD Thesis, University of Edinburgh 1993.
- B. Blumberg. *Old Tricks, New Dogs: Ethology and Interactive Creatures*. PhD Thesis, Massachusetts Institute of Technology, 1997.

In addition to this, it can be appreciated that aspects of the approach described here in respect of simulation of creatures can be employed more generally to simulation. In particular, other situations may be considered in which a simulation may be split into a first part which defines states and the possible transitions between states - a state machine - and two or more alternative second parts which are adapted to implement these transitions at different levels of complexity.

A general "first part" is shown in Figure 3A, and closely resembles the state machine of Figure 2. Three states are shown (A, B and C), together with possible transitions between them, and (in this case), probabilities for each transition. The general "second part", or collection of second parts, is shown in Figure 3B. It is desirable first to have a rule evaluated with reference to at least one condition determined by the computational system. This may be, for example, processor time available to carry out the simulation or predicted to be available over a certain period. This may apply where the simulation is a single simulation rather than one of a group, and that the constraint on the simulation is that it must be carried out within a fixed period of time

despite having varying computational resources available. The condition may also be an indication of the significance of the simulation (rather as in the case of animats - there may be some determination as to which simulated entities are important, or important at that time, and computational resources may be preferentially allocated to such entities - note here, however, that the "simulated entities" may be simply processes rather than specifically creatures).

10

On the basis of the evaluation of the rule, alternative second portion choices 503, 504, 505 are followed for evaluating each of the transitions in Figure 3A (as is shown here, there may be more than two alternative levels of complexity available, rather than simply two). Each of these second choices has the same set of transitions, but different rules or methods for implementing these rules such that the choices are of different levels of computational complexity. Here, for example, choice 503 simply provides transition probabilities to be evaluated over a particular period of time. Choice 504 involves further variables a, b and c in calculating transitions - this will be of greater computational complexity than for choice 503, but will, preferably, be such that choice 503 provides an approximation to the process of choice 504 (so that both can reasonably be regarded as implementations of the first portion shown in Figure 3A). Choice 505 involves further variables d, e and f, and is thus presumed for this example only to be of still greater computational complexity - though it is appreciated that the presence of additional variables will not necessarily result in computational complexity, and that this may also result from other causes (such as functions that are more

complex to evaluate). These different choices are used for simulation - Figure 3B shows a loop structure with periodic re-evaluation of the rule 502 after a waiting step 506 (suggesting that the simulation can change levels of complexity during the process - which may be desirable if, for example, the computational resources available vary), but in other cases the simulation may simply continue until an end condition is reached.

10 As will be appreciated from the above, such an approach to simulation is not limited in its application to large-scale simulation of multiple creatures. The approach can be applied to simulations of single creatures, entities or even processes with different levels of computational
15 resource available, or may be applied to simulations of plural entities or processes.

The reader's attention is directed to all papers and documents which are filed concurrently with or previous to
20 this specification in connection with this application and which are open to public inspection with this specification, and the contents of all such papers and documents are incorporated herein by reference.

25 All of the features disclosed in this specification (including any accompanying claims, abstract and drawings), and/or all of the steps of any method or process so disclosed, may be combined in any combination, except combinations where at least some of such features
30 and/or steps are mutually exclusive.

Each feature disclosed in this specification (including any accompanying claims, abstract and drawings), may be

replaced by alternative features serving the same, equivalent or similar purpose, unless expressly stated otherwise. Thus, unless expressly stated otherwise, each feature disclosed is one example only of a generic series
5 of equivalent or similar features.

The invention is not restricted to the details of the foregoing embodiment(s). The invention extends to any novel one, or any novel combination, of the features
10 disclosed in this specification (including any accompanying claims, abstract and drawings), or to any novel one, or any novel combination, of the steps of any method or process so disclosed.

CLAIMS:

1. A method of simulating a creature for use in two
5 different complexities of simulation, the method
comprising utilising a model of the creature that
comprises at least two portions,
a first portion which contains functions for use in both
of said different complexities of simulation; and
10 a second portion comprising two alternative versions, a
first version for use in one of said different
complexities of simulation, and a second version for use
in the other of said different complexities of simulation.
- 15 2. A method as claimed in claim 1, wherein said first
portion comprises a behaviour selection mechanism arranged
to select the behaviour of said creature.
3. A method as claimed in claim 2, wherein said behaviour
20 selection mechanism is arranged to select the behaviour
based upon at least one of: the current behavioural state;
one or more internal state variables of the creature; the
environment surrounding the creature; one or more sensory
inputs to said creature.
- 25 4. A method as claimed in claim 2 or claim 3, wherein
said behaviour selection mechanism consists of a set of
mutually exclusive behavioural states.
- 30 5. A method as claimed in any one of the above claims,
wherein the second version is for use in the less complex
of the simulations, and is arranged to approximate the
functionality of the first version.

6. A method as claimed in any one of the above claims, wherein the first version utilises a neural network.

5 7. A method as claimed in claim 2, or any claim dependent thereto, wherein said second portion is arranged to execute the selected behaviour.

8. A method as claimed in any one of the above claims,
10 wherein the first version utilises a three dimensional physical simulation of the animat, and the second version utilises a parameterised model of the animat to approximate movement.

15 9. A method of simulating the activities of a plurality of creatures, the method comprising utilising at least two modes of simulation: a first mode arranged to simulate the activities of all of said creatures; and a second mode arranged to simulate the activity of at least one of said
20 creatures at a more detailed level than said first mode; wherein the model of a creature simulated in both modes of simulation comprises at least two portions:
a first portion which contains functions arranged for use in both of said modes of simulation; and
25 a second portion comprising two alternative versions, a first version for use in said first mode of simulation, and a second version for use in the second mode.

10. A method of simulating a process at two different
30 levels of complexity, the method comprising utilising a model that comprises at least two portions,
a first portion which contains functions for use in both of said different complexities of simulation; and

a second portion comprising two alternative versions, a first version for use in one of said different complexities of simulation, and a second version for use in the other of said different complexities of simulation.

5

11. A method as claimed in claim 10, further comprising evaluating one or more conditions to determine a result of a rule for selecting which of the two alternative versions of the second portion to use in simulating the process.

10

12. A method as claimed in claim 10 or claim 11, wherein the second version is for use in the less complex of the simulations, and is arranged to approximate the functionality of the first version.

15

13. A method as claimed in any of claims 10 to 12, wherein the first version utilises a neural network.

20

14. A computer program arranged to program a processor of suitable computing apparatus to perform the method as claimed in any one of the above claims.

15. A recordable medium comprising a computer program as claimed in claim 14.

25

16. A simulator device arranged to simulate a creature in two different complexities of simulation; the device being arranged to utilise a model of the creature that comprises at least two portions;

30

a first portion which contains functions used in both of said different complexities of simulation; and a second portion comprising two alternative versions, a first version used in one of said different complexities

of simulation, and second version used in the other of said different complexities of simulation.

ABSTRACTSIMULATION AT TWO DIFFERENT LEVELS OF COMPLEXITY

5

A technique for consistently simulating at two different complexities of simulation is described. This technique is particularly relevant to simulation of a creature-like autonomous agent (biological or artificial or hybrid).

10 The model, in the particular case a model of the creature, comprises at least two portions. The first portion contains functions used in both of the different complexities of simulation. The second portion comprises two alternative versions, a first version for use in one
15 complexity of simulation, and a second version for use in the other complexity of simulation.

[Figure 1]

20

25




FIGURE 1




FIGURE 2




FIGURE 3A

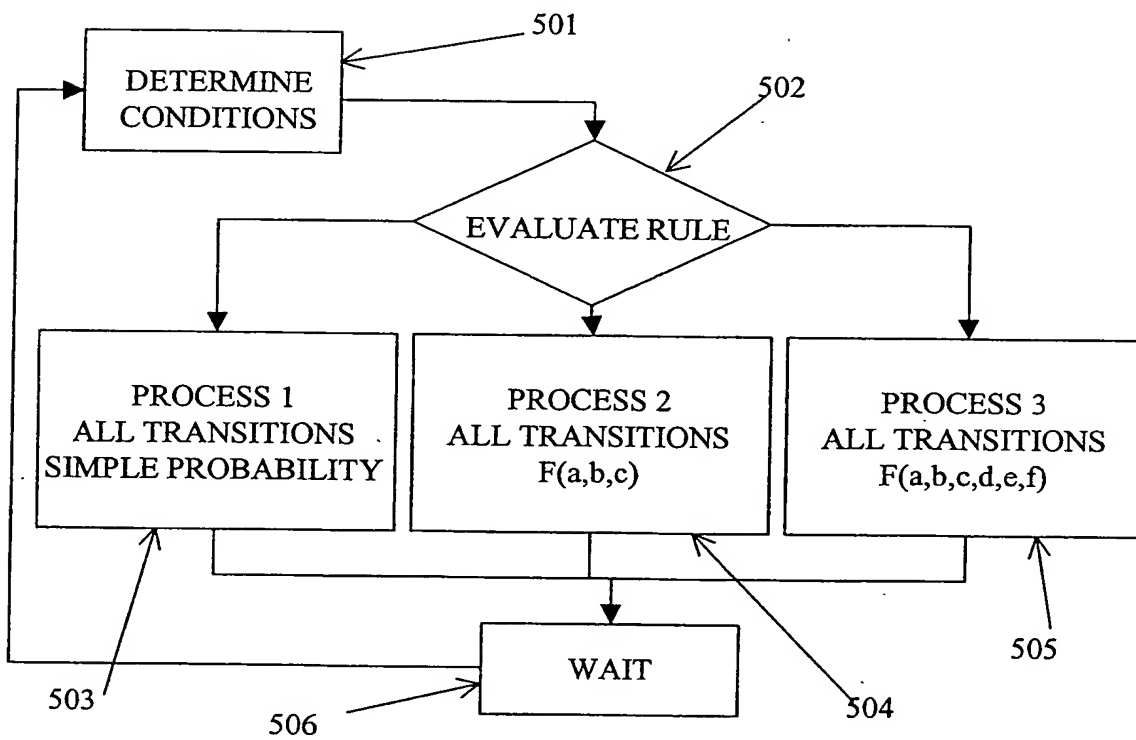


FIGURE 3B

